Journal of Machine Learning Research 1 (2016) 00-00 Submitted 4/00; Published 10/00

Sources identification using shifted non-negative matrix factorization combined with semi-supervised clustering

Filip L. Iliev FILIEV@UCSC.EDU
Physics and Chemistry of Materials Group
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM, USA

Valentin G. Stanev VSTANEV@UMD.EDU
Physics and Chemistry of Materials Group
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM, USA

Velimir V. Vesselinov VVV@LANL.GOV
Computational Earth Science Group
Earth and Environmental Sciences Division
Los Alamos National Laboratory
Los Alamos, NM, USA

Boian S. Alexandrov BOIAN@LANL.GOV
Physics and Chemistry of Materials Group
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM, USA

Editor: Dr. Daniel Lee

arXiv:1612.03950v1 [cs.LG] 12 Dec 2016

Abstract

Non-negative matrix factorization (NMF) is a well-known unsupervised learning method that has been successfully used for blind source separation of non-negative additive signals. NMF method requires the number of the original sources to be known *a priori*. Recently, we reported a method, we called NMFk, which by coupling the original NMF multiplicative algorithm with a custom semi-supervised clustering allows us to estimate the number of the sources based on the robustness of the reconstructed solutions. Here, an extension of NMFk is developed, called ShiftNMFk, which by combining NMFk with previously formulated ShiftNMF algorithm, Akaike Information Criterion (AIC), and a custom procedure for estimating the source locations is capable of identifying: (a) the number of the unknown sources, (b) the eventual delays in the signal propagation, (c) the locations of the sources, and (d) the speed of propagation of each of the signals in the medium. Our new method is a natural extension of NMFk that can be used for sources identification based only on observational data. We demonstrate how our novel method identifies the components of

©2016 F.Iliev, V.Stanev, V.V.Vesselinov and B.S.Alexandrov.

synthetic data sets, discuss its limitations, and present a Julia language implementation of ShiftNMFk algorithm.

Keywords: Non-negative matrix factorization; Semi-supervised learning; Source identification; Sources allocation; Delay propagation

1. Introduction

Nowadays large datasets are collected at various scales, from laboratory to planetary, (Kitchin and McArdle, 2016; Chen et al., 2012). The collected data are often high-dimensional and contains many state variables. For example, large amounts of data are gathered by distributed sets of asynchronous sensor arrays that, via remote sensing, can measure considerable number of parameters, such as: temperature, sounds, stress, tension, power, flow, vibrations, RF and IR spectra, satellite images, and many others. Usually, each record of a sensor in the array is a mixture of signals originating from unknown number of physical sources with undetermined locations and strengths. The analysis of such types of data, especially related to threat reduction, nuclear nonproliferation, and environmental safety, could be critical for emergency responses. One of the main goals of such analysis is to identify and localize set of unknown number of events and activities, such as, seismic events, contamination sources, pollution emissions, climate change impacts, etc. Currently, such type of analysis requires complex, poorly constrained and uncertain physics-based inverse-model methods where complex, high-resolution, computationally-intensive numerical models are typically be needed to simulate the governing process related to signal propagation through the media of interest. An alternative approach is the model-free analysis based on the Blind Source Separation (BSS) techniques (Belouchrani et al., 1997).

In the typical BSS problem, the observational data, V, ($V \in M_{n,m}(R)$), is formed by a linear mixing of d unknown original signals, H, ($H \in M_{d,m}(R)$), blended by an also unknown mixing matrix, W, ($W \in M_{n,d}(R)$), i.e.,

$$V_{n,m} = \sum_{d} W_{n,d} H_{d,m} + \epsilon, \tag{1}$$

where ϵ is a vector ($\epsilon \in V_m(\mathbb{R})$) and denotes presence of possible noise or unbiased errors in the measurements (also unknown). If the problem is solved in a temporally discretized framework, the goal of the BSS algorithm is to retrieve the d original signals, H, that have produced n observational mixtures of these signals, V, recorded by a set of sensors. Here n is the number of the sensors, d is the number of the unknown signals (sources) observed in the collected data (V), and m is the number of discretized moments in time at which the signals are recorded by the sensors. Note that the algorithm expects all the observations to be at the same times. The temporal spacing between does not need to be uniform. If the data are not collected at the same times, interpolation techniques can be applied to pre-process the data.

Since both factors H and W are unknown (the size d of these matrices is also unknown, because we do not know how many original signals/sources have been mixed in sensors' records), the main difficulty in solving any BSS problem is that it is under-determined.

There are two widely-used approaches established to resolve the BSS underdetermination: Independent Component Analysis (ICA), (Herault and Jutten, 1986; Amari et al.,

1996), and Non-negative Matrix Factorization (NMF), (Paatero and Tapper, 1994; Lee and Seung, 1999). The main idea behind ICA is that, while the probability distribution of a linear mixture of the sources in V is expected to be close to a Gaussian (according to the Central Limit Theorem), the probability distribution of the original sources could be non-Gaussian. Hence, ICA aims to maximize the non-Gaussian characteristics of the estimated sources in H, with the goal to find the maximum number of statistically independent original sources that, when mixed by W, reproduce the observational data (Eq.1). The second approach, NMF, is a well-known unsupervised learning method, created for parts-based representation, (Fischler and Elschlager, 1973) in the field of image recognition, (Paatero and Tapper, 1994; Lee and Seung, 1999), and successfully leveraged for Blind Source Separation, that is, for decomposition of mixtures formed by various types of non-negative signals, (Cichocki et al., 2009). In contrast to ICA, NMF does not impose statistical independence or any other constrain on statistical properties (i.e., NMF allows the original sources to be correlated); instead, NMF enforces a non-negativity constraint on the original signals in H and their mixing components in W. NMF can successfully decompose large sets of non-negative observations, H, by leveraging the multiplicative update algorithm introduced by Lee and Seung (Lee and Seung, 1999). However, NMF requires *a priori* knowledge of the number of the original sources.

Recently we reported a method, we called NMFk, which by coupling the original multiplicative algorithm with a custom semi-supervised clustering allows to identify the number of the unknown sources based on the robustness and reproducibility of the solutions (Alexandrov and Vesselinov, 2014). Here, we report an extension of NMFk, called ShiftNMFk, which by combining NMFk with a) previously developed ShiftNMF algorithm, (Mørup et al., 2007), b) Akaike Information criterion (Akaike, 2011), and c) with a custom source-location procedure, is capable of evaluation the number of the original sources and estimation of their eventual delays, locations, and speed in the medium. This new method, ShiftNMFk, is a natural extension of NMFk needed for the identification of the unknown number of sources and their locations when the signals originating from these sources have a finite speed of propagation in the considered medium.

2. Methodology

2.1 NMF algorithm

Mathematically NMF is presented by Eq. 1, with the observations being the rows of $V_{n,m}$. The algorithm returns the decomposition through the mixing matrix $W_{n,d}$ and source matrix $H_{d,m}$ with ϵ being the residual noise. The rows in V correspond to the number of sensors while the rows of H correspond to the number of sources. The number of sensors has to be greater than the number of sources. For NMF to work, the problem must be amenable to a non-negativity constraint on the sources H and mixing matrix W. This constrain leads to reconstruction of the observations (the rows of matrix V) as linear combinations of the elements of H and W that cannot cancel mutually.

The NMF algorithm starts with a random guess for H and W, and proceeds by minimizing the cost (objective) function, O, which in our case is the Frobenius norm,

$$O = \frac{1}{2}\|V - W * H\|_F^2 = \frac{1}{2}\sum_{n,m}\left(V_{n,m} - \sum_d W_{n,d}H_{d,m}\right)^2 \quad (2)$$

during each iteration. Minimizing Frobenius norm, is equivalent to representing the discrepancies between the observations, V, and the reconstruction, $W * H$, as white noise. In order to minimize O, we use the gradient descent approach based on the NMF multiplicative updates proposed by Lee and Seung (Lee and Seung, 1999). During each iteration, the algorithm first minimizes O by holding W constant and updating H, and then holds H constant while updating W, (Eq. 3 and 4). It is proven that the norm, (Eq. 2), is non-increasing under these update rules and invariant when an accurate reconstruction of V is achieved, (Lee and Seung, 1999),

$$H_{d,m} = H_{d,m}\frac{(W^T V)_{d,m}}{(W^T W H)_{d,m} + \epsilon}, \quad (3)$$

$$W_{n,d} = W_{n,d}\frac{(V H^T)_{n,d}}{(W H H^T)_{n,d} + \epsilon}, \quad (4)$$

where the small constant (e.g., $\epsilon \sim 10^{-9} - 10^{-16}$) was added to the denominator to avoid division by zero.

2.2 NMFk algorithm

While powerful and producing results easy for interpretation, the classical NMF requires *a priory* knowledge of the number of the original sources. By coupling the NMF method with a custom semi-supervised clustering we demonstrated that the number of the original sources can be estimated based on the robustness and reproducibility of the NMF solutions. This approach was leveraged to decompose the largest available dataset of human cancer genomes (Alexandrov et al., 2013a) and then extended for decomposition of physical signals/transients (Alexandrov and Vesselinov, 2014). The method, called NMFk, explores consecutively the range of possible numbers of original sources, from $1, 2..., n$, and then estimates the accuracy and reproducibility of solutions with different number of original sources. Note that n is the number of sensors; in general, we cannot identify more sources than the number of sensors without additional information. This means that NMFk performs n sets of independent runs using a different (fixed in each run) number of original sources, $d = 1, 2, ..., n$. Each run performs a large number, M, of NMF simulations, with random initial conditions, and, therefore, after each run we obtain a set U_d of M solutions, each with two matrices, H_d^i and W_d^i, (for d original sources, and $i = 1, 2, ..., M$):

$$U_d = ([H_d^1; W_d^1], [W_d^2; W_d^2], ..., [H_d^M; W_d^M]) \quad (5)$$

Next, NMFk uses a custom semi-supervised clustering to assign each of these M solutions (in U_d for each d) to one of the specific clusters. This custom semi-supervised method is a k-means clustering that keeps the number of solutions in each cluster equal. For example,

after a run with $M = 1,000$ simulations (performed with random initial conditions), each one of the d identified clusters, representing the unknown d sources, has to contain exactly $1,000$ solutions.

Note that we have to enforce the condition that each of these d clusters contains equal number of solutions, since each NMF-simulation contributes only one possible solution for each source. During clustering the similarity between sources is measured using the cosine distance (Pang-Ning et al., 2006).

The main idea for estimating the unknown number of sources is to use the separation of the clusters as a measure of how good is a particular choice for the number of sources, d. On the under-fitting side, for solutions with d less than the actual number of sources, we expect that the clustering could be good because several of the actual sources could be combined to produce one "super-cluster". However, clustering will break down significantly with the over-fitting, when d exceeds the true number of sources since even if the average reconstruction error of the solution is small, we do not expect the solutions to be well clustered (since there is no unique way to reconstruct the solutions with more than the actual number of sources, and at least some of the clusters will be artificial, rather than real entities). Thus, if we estimate the degree of clustering for the solutions with different number of sources, and plot it as a function of d, we expect a sharp drop after we exceed the actual number of sources (Alexandrov and Vesselinov, 2014). After the clustering, the average Silhouette width (Rousseeuw, 1987) is computed to estimate how well the solutions are clustered (Alexandrov and Vesselinov, 2014). Using the clusters based on the clustering of the solutions (i.e., clustering of the signals matrix and obtaining the corresponding clusters of the mixing matrices, W_d), the average solutions, i.e., the centroids of these clusters, are computed. The optimal number of sources, r, is picked by selecting r to be the number of sources that demonstrate both (a) an acceptable reconstruction of the observation matrix V and (b) a high average Silhouette width (i.e., close to one). Here in addition of these two metrics to select the unknown number of sources, we further extend NMFk methodology by adding the a new metric based on the Akaike Information Criterion (AIC) which provides evaluation of solution parsimony where solutions with smaller number of estimated parameters are preferred if they perform similarly to solutions with larger number of estimated parameters (Akaike, 2011).

2.3 ShiftNMF algorithm

Serious limitation of the standard NMF, as well as the NMFk, approach is the assumption of infinite propagation speed of the signals. Practically all the sources are assumed to be mixed instantaneously at the sensors. Thus, a natural extension of the NMFk method is to take into account the delays in the signal propagation from the sources to the sensors (Smaragdis and Brown, 2003), caused by the different positions of the physical sources in space, combined with the finite speed of propagation of the signals into the considered medium. One way of treating this problem is to write it as convolutive NMF (Cichocki et al., 2009), delays of the signals being a particular form of convolution:

$$V_{n,m} = \sum_{p} \sum_{d} W^{p}_{n,d} H^{p}_{d,m} + \epsilon, \tag{6}$$

where p is a time-shift operator. Unfortunately, solving convolitive NMF, usually presents a formidable numerical problem. Here we use an alternative approach, and leverage previously developed ShiftNMF algorithm, (Mørup et al., 2007), designed specifically for finding signals with delays. This method introduces an additional minimization procedure (see below) needed to estimate the locations of the original physical sources, whose delayed signals are forming the mixtures recorded by the sensors. Below we describe the key features of the ShiftNMF algorithm, following closely the original work by Mørrup, Madsen and Hansen (Mørup et al., 2007).

The mathematical representation of the ShiftNMF algorithm is very similar to the classic NMF but features an additional matrix τ which incorporates the potential delays (time shifts) in observing the same signal at different sensors.

The $\tau_{n,d}$ matrix element denotes the delay from the d^{th} original signal to the n^{th} sensor, and using it we can write

$$V_{n,m} = \sum_d W_{n,d} H_{d,m-\tau_{n,d}} + \epsilon, \tag{7}$$

With delays, just as with more general forms of convolution, it is advantageous to use the Fourier space, in which the previous equation becomes

$$\widetilde{V}_{n,m} = \sum_d W_{n,d} \widetilde{H}_{d,f} e^{-i2\pi \frac{f-1}{M} \tau_{n,d}} + \widetilde{\epsilon}_f, \tag{8}$$

where the $\widetilde{\cdot}$ denotes the image in Fourier space. Somewhat more succinct version of this equation is

$$\widetilde{V}_f = \widetilde{W}^{(f)} \widetilde{H}_f + \widetilde{\epsilon}_f; \text{ where } \widetilde{W}^{(f)} = W_{n,d} e^{-i2\pi \frac{f-1}{M} \tau_{n,d}}, \tag{9}$$

and note that we have used the elementary property of the Discrete Fourier Transform (DFT) to convert time shifts into phase factors.

As we can see, ShiftNMF returns not only the source matrix, H, and the mixing matrix, W, but also an additional $n \times d$ mapping matrix, τ, that contains the delays (as integer numbers) of the signal from each original source, d, to each sensor, n. Mørrup, Madsen and Hansen proposed an algorithm, which they called ShiftNMF, to find simultaneously W, H and τ. ShiftNMF uses a similar strategy of multiplicative updates as the conventional NMF, but jumps into the frequency domain and back at each update. In Fourier space the nonlinear shift mapping becomes a family of DFT transformed H matrices with the shift amount taken by the unknown τ matrix. Thus the delayed version of the source signal, to the n^{th} channel, is,

$$\widetilde{H}^{(n)}_{d,f} = \widetilde{H}_{d,f} e^{-i2\pi \frac{f-1}{M} \tau_{n,d}}, \tag{10}$$

and the Frobenius norm that has to be minimized is,

$$O = \frac{1}{2} \sum_{n,m} \left(V_{n,m} - \sum_d W_{n,d} H_{d,m-\tau_{n,d}} \right)^2 = \frac{1}{2M} \left\| \widetilde{V_f} - \widetilde{W}^{(f)} \widetilde{H_f} \right\|_F^2, \tag{11}$$

where the last equality holds because of the Parseval's identity. It is clear that ShiftNMF has to update three matrices, H, W, and τ.

The updates of the mixing matrix, W, is done in the same way as the conventional NMF but incorporating one of the changed (by τ) H^n matrix (which is also non-negative).

$$W_{n,d} = W_{n,d} \frac{V_{n,:} H_{d,:}^{(n)^T}}{W_{n,:} H^{(n)} H_{d,:}^{(n)^T}} \tag{12}$$

The updates of the H matrix are multiplicative, given by the gradient of the Frobenius norm, O, in Fourier space,

$$\widetilde{G}_f = \frac{\partial O}{\partial \tilde{H}_f} = -\frac{1}{M} \tilde{W}^{(f)^H} (\tilde{V}_f - \tilde{W}^{(f)} \tilde{H}_f) \tag{13}$$

By taking the inverse Fourier of the gradient, $\widetilde{G}_f$, and separation into its positive and negative parts in the form, $\widetilde{G}_f = \widetilde{G_{d,n}^+} + \widetilde{G_{d,n}^-}$, the corresponding gradient, $G_{d,n} = G_{d,n}^+ + G_{d,n}^-$, in the time domain can be derived and used to update H,

$$H_{d,m} = H_{d,m} \left(\frac{G_{d,n}^-}{G_{d,n}^+} \right)^{\alpha} \tag{14}$$

Here the time shift is already incorporated into the matrix W and the gradient, while α is a tunable time-step factor ($\alpha \to 0$, means very small steps in the negative gradient direction).

Because the delays are unconstrained, the shift matrix τ is estimated by Newton-Raphson method which simply looks for the minimum of a function with the help of its gradient and Hessian,

$$T \leftarrow T - \eta B^{-1} g \tag{15}$$

Here T is vectorized form of the matrix τ (with dimensions $(n \times d, 1)$), and g is the gradient of the Frobenius norm, O, with respect to the matrix τ. Within O, the delay is once again folded into the mixing matrix, W. In Eq.15, B is the Hessian of the same Frobenius norm, O, while η is a tunable constant that can be changed within the algorithm to ensure faster convergence. Alone, however, this update procedure is sensitive to local minima. To mitigate that τ is updated during every 20^{th} iteration, using a specific cross-correlation procedure (Mørup et al., 2007), which has the effect of "kicking" the ShiftNMF solution out of the local minimum it may have become trapped into.

It is important to note that this algorithm determines the relative delays (the time shifts of the same signal arriving at different sensors). Thus τ is determined up to a constant – a global time shift of all delays would only lead to rearrangement of the matrices in Eq.7. The algorithm resolves this ambiguity by centering the delays at zero, and splitting them in positive and negative values.

2.4 ShiftNMFk algorithm

The general idea in ShiftNMFk is to substitute the NMF-simulations in each run of our NMFk method with ShiftNMF-simulations. As outlined above in Section 2.2, to estimate the unknown number of original sources, we have to perform n sets of ShiftNMF-simulations, which we call runs, and the simulations of each of these runs uses same number, d, of

original sources, $d = 1, 2, ..., n$. Thus, each run performs a large number, M, of ShiftNMF simulations (with random initial conditions) with the same number of sources. With each run we result with a set of solutions, U_d, containing M solutions, each with three matrices, H_d^i, W_d^i, and τ_d^i (with d original sources),

$$U_d = ([W_d^1; W_d^1; \tau_d^1], [H_d^2; W_d^2; \tau_d^2], ..., [W_d^M; W_d^M; \tau_d^M]) \tag{16}$$

Finally, we apply our custom semi-supervised clustering to each of these sets, U_d, to extract the average solutions (and their delays), and to estimate their robustness for each d, by comparing the average Silhouette width of the clusters and average reconstruction error, as described above, in NMFk.

However, performing the ShiftNMF algorithm, many times and with random initial conditions, we found that during the simulations (in each run, U_d) shiftNMF may converge to different (often very dissimilar) solutions, while trying to minimize the norm, O, and frequently may stop before reaching a good reconstruction. This is due to several factors, such as specific initial conditions, the ratio between the number of sources and number of sensors, specific shape of the signals and delays, etc. The accurate reconstruction, for example, depends on the level of correlation between the waveforms of original signals (see, below, Figure 1).

Thus, in a sizable percentage of the ShiftNMF simulations in each run, i.e., for each choice of the number of the original sources, d, the algorithm exits (many times) with a poor reconstruction of the observational matrix, V.

To overcome these limitations, we employ two elimination criteria needed to derive (from a reasonably sized pool of M ShiftNMF simulations) in each run robust solutions that reproduce accurately the observations and are physically meaningful.

2.4.1 ShiftNMFk elimination criteria

In ShiftNMFk, we use the following two elimination criteria to discard (a) solution outliers that do not reconstruct sufficiently well the observation matrix, V, and (b) solutions that do not satisfy a general physical condition. The physical condition we are leveraging is a constraint related to existence of a maximum signal delay (between all the sensors) that arises from the finite size of the sensor array and finite speed of propagation of the signals in the medium. Let us consider in details these two selection criteria.

1. *Remove the outliers*: We discard the ShiftNMF solutions that fail to minimize the discrepancy between the observational matrix, V, and its reconstruction, $W * H$, i.e., the solutions that fail to minimize the norm in (Eq.2) to a reasonable degree. Specifically, we throw away solutions with a ratio of Frobenius norms, $R = \frac{\|V - W*H\|_F}{\|V\|_F}$ larger then a certain value. In the analysis presented below, we used $R < 10\%$. This criteria removes ShiftNMF solutions that are crude representations of the observation matrix, V, and can be considered as outliers.

2. *Impose a maximum time delay in the array*: Our second criteria for elimination of ShiftNMF solutions focuses on exploring the signal delays. In a specific sensor array, the sensors are separated at various distances. Among these distances, there is a

largest one, which we will call the largest distance in the array, L_{max}. Next, a maximum time for travel through the array, t_{max}, can be defined as the largest distance in the array divided by the speed of propagation: $t_{max} = L_{max}/v_{max}$. Furthermore, for a given signal (and speed of propagation in the considered medium, v) none of the differences between any two of the signal's delays (to any two of the sensors in the array) could be larger than the time needed for the signal to propagate through the maximum distance, L_{max}. Based on that, we can apply a physical constraint like this: The spread of the delays of each of the signals over all of the sensors is limited by the maximum travel time, t_{max}. To implement this criterion, we can apply the information about the physical coordinates of the sensors in the array (typically this information is known) and the speeds of propagation the signals (could be unknown). However, this criterion can be also implemented without a detailed knowledge of the metrics of the sensor array. This can be done by evaluation of the standard deviation, $std(\tau_{j,i})$, and average value of the delays, $mean(\tau_{j,i})$, from each simulation (in the specific run) associated with a given signal j. We apply the coefficient of variation, $CV(\tau_{j,i})$ of the delays from all the sources to all the sensors, $CV(\tau_{j,i}) = std(\tau_{j,i})/mean(\tau_{j,i})$. If the spread of the delays of a given signal over all the sensors (in the simulations in a given run), CV, is greater than a certain cut-off, e.g., $CV \geqslant 0.8$, the corresponding solution can be considered unphysical, and rejected.

The above two conditions are defined to be sufficiently general and expected to be needed for any type of problem. However, if need they may be adjusted when specific real problems are solved without any lost of generality of our algorithm.

2.4.2 ShiftNMFk semi-supervised clustering

After leveraging the above elimination criteria and disregarding some of the shiftNMF solutions, we result with a pool of solutions that can be used in NMFk semi-supervised clustering. We cluster the solutions from this pool, corresponding to different number of original sources, d, and following the strategy of NMFk compare the average Silhouette width for each number of original sources and the average value of the corresponding reconstruction error, O, for each number of original sources d. Here the average Silhouette width evaluates the robustness of the solutions with given number of original sources, d, while the reconstruction error, O, evaluates how accurate these average solutions reconstruct the observation matrix, V. Finally, the optimal number of original sources is determined by finding where a reasonable average reconstruction error is achieved for a maximum average Silhouette width. To achieve this and to avoid over-fitting, we make use of the Akaike Information criterion (Akaike, 2011). First, we take all possible numbers of sources that lead to an average Silhouette width above certain threshold (for example, 0.7). Then, we select the value of d that minimizes

$$AIC = 2d - 2\ln(L) = 2d + N\ln\left(\frac{O^{(d)}}{N}\right), \tag{17}$$

where L is the likelihood of the model with given number of sources d. Note that, $N = n*m$, is the total number of data points (m is the number of the time points in the records). We

define L using the average reconstruction error $O^{(d)}$ of the solutions we have kept in the clustering procedure for this particular number of sources d: $\ln(L) = -(N/2)\ln(O^{(d)}/N)$.

The AIC is a standard measure of the relative quality of statistical models, which takes into account both the likelihood (depending on the average reconstruction error) and the degrees of freedom (depending on the number of sources) needed to achieve this likelihood. AIC penalizes models with higher number of sources. If we compare AIC estimates for 2 solutions with the similar likelihoods (average reconstruction error) but different number of sources, the AIC will be lower for the model with smaller number of sources and this solution should be preferred.

2.4.3 Retrieving original source locations

Having completed the ShiftNMFk procedure, we now have the needed information to estimate the locations of the original sources and the speeds of propagation of the signals. For this purpose, we use a weighted Least-squared minimization procedure to estimate the coordinates of the sources based on coordinates of the sensors and the signal delays, where if,

$$r_{j,i} = \sqrt{(x_j^s - x_i^d)^2 + (y_j^s - y_i^d)^2} \tag{18}$$

the minimization function F is defined like this:

$$F = \sum_{j}^{n} \sum_{i \neq i^*}^{d} \left(\frac{(\tau_{j,i^*} - \tau_{j,i}) - \frac{(r_{j,i^*} - r_{j,i})}{v_j}}{\sqrt{\sigma_{j,i^*}^2 + \sigma_{j,i}^2}} \right)^2 \tag{19}$$

Here i^* denotes the closest sensor to the source with number j (the sensor with the smallest delay for jth source), and d and n are the number of sources and sensors, respectively. For simplicity, we will assume that the speed of propagation of each of the signals, v_j, has the same value, v (but the method is not restricted to this case). Also, note that we have assumed two-dimensional space, but F in its general form can be used in arbitrary dimension, provided $r_{j,i}$ is suitably modified. Further, the equation for $r_{i,j}$ gives the distance from a source j to sensor i, and x_j^s and y_j^s are the coordinates of source j, while x_i^d and y_i^d are the coordinates of sensor i. The standard deviations, $\sigma_{j,i}$, are the sample standard deviations of $\tau_{j,i}$ for a fixed source j and sensor i in the run. The samples are obtained from ShiftNMF simulations (in each run and after the elimination criteria), and we assume that all the $\tau_{j,i}^p$, $p = 1, 2, ..., M^*$ in the ShiftNMF simulations in any of the runs are normally distributed.

The coordinates of all the sensors are known. The coordinates of the sources are the unknown variables, which, along with the speeds of signals propagation, can be found by minimizing the objective function F. Indeed, the difference between the signal's delays to two (arbitrary) sensors, i_1 and i_2, should be the same as the difference between the corresponding distance $r_{1,2}$ between these sensors divided by the speed of propagation. Hence, by minimizing F, we are trying to find the coordinates of the sources and speeds that makes this true. Minimizing F also allows us to propagate the uncertainty/error of our ShiftNMFk estimates into our source location estimates by incorporating the sample standard deviations of $\tau_{j,i}$ for each source.

Practically, the minimization of F is performed using the NLopt.jl Julia package and running the nonlinear minimization ∼1,000 times with different initial (random) conditions and suitable constraints (see below).

After the minimization we get rid of the outliers using two different criteria: A) we only keep the best 50% of the solutions defined by how well the function F was minimized, and B) From these solutions we set the median position coordinates as the likely coordinates of our sources and we get rid of the 50% of solutions whose coordinates are far away from these likely coordinates. In this way, for each source, we left with relatively tightly clustered solutions around the likely coordinates.

2.4.4 Uncertainties of the ShiftNMFk estimates

We use Bayesian analysis leveraged with Markov Chain Monte Carlo (MCMC) sampling to obtain the posterior probability distribution functions (PDFs) of the estimated, by the nonlinear optimization, speed and coordinates of the sources. The PDFs are obtained based on the estimated delays and following Bayes theorem (Weiss, 1996; Oakley and O'Hagan, 2004), using a likelihood function defined as $e^{(-\chi^2/2)}$. The analysis was performed using the Robust Adaptive Metropolis Markov Chain Monte Carlo (MCMC) algorithm (Vihola, 2012), implemented in the open-source computational framework MADS (Model Analysis and Decision Support: http//mads.lanl.gov). Using the results of this Bayesian analysis, we can define a region of likelihood where the position of the original sources can be found with a given probability.

3. Results

3.1 Synthetic datasets

Our synthetic datasets are constructed by generating various observation matrices using a semi-random approach. Specifically, we start with several (two, three, and four) basic waveforms, as original signals, H, which we then randomly mix and shift, with respect to each sensor, by randomly generating the mixing matrix, W, and delay matrix, τ. We are using sensor arrays with 16, 18, or 24 sensors distributed on rectangular discrete lattices. In this way, for each choice of the original sources, H, mixing matrix, W, or delay matrix τ, we obtain a different observation matrix, V. We explored the success rate of ShiftNMFk on a large number of observation matrices generated in this way.

3.2 Impact of the number of ShiftNMF iterations

We explored the optimal number of iterations, I, for each ShiftNMF simulation, that allows a reasonable average reconstruction error. Our results demonstrate that after a certain number, I_{max}, increasing further the number of the iterations does not lead to any improvement of the final results. This is because often the algorithm stopped by its internal convergence criteria, before reaching the maximum number of iterations. We found that, $I_{max} = 50,000$, produces acceptable results.

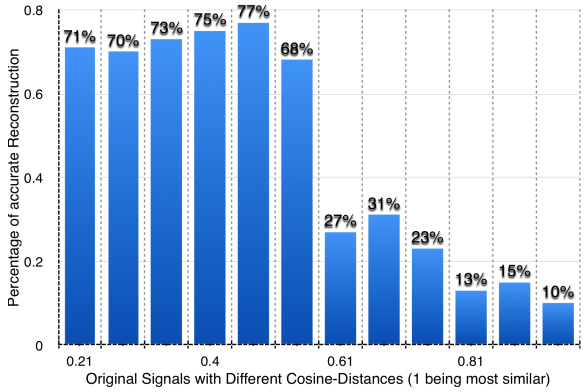

Figure 1: Percentage of correct reconstruction obtained by the original ShiftNMF algorithm, for 3 source signals with different levels of cosine similarity.

3.3 Impact of signal correlations on ShiftNMF results

To be able to predict the number of the original sources, a proper understanding of the limitations of the ShiftNMF simulations with random initial conditions is required. We performed a huge number of ShiftNMF simulations to unravel these limitations. Similar to the conclusions about classical NMF performance in Ref. (Alexandrov et al., 2013b), ShiftNMF works better if the original signals are not strongly correlated. The reason is easy to understand: if two of the original signals are strongly correlated, the algorithm can easily assume that the mixtures recorded by the sensors are caused by one source instead of two nearly identical sources, returning a wrong original signal that, however, provides a relatively decent reconstruction.

Changing the correlation between two of three initially generated signals, we studied the recognition success rate by comparing the cosine distance between the generated original signals and these derived by the shiftNMF algorithm (one simulation). Equation 20 describes the cosine distance comparing vectors a and b where a_i and b_i are individual elements of each vector. For a successful recognition, we accept the cosine distance, $\rho(a,b) \geqslant 0.95$, where

$$\rho(a,b) = 1 - \frac{\sum_{i=1}^{n} a_i b_i}{\sqrt{\sum_{i=1}^{n} a_i^2}\sqrt{\sum_{i=1}^{n} b_i^2}} \tag{20}$$

Figure 1 presents the relationship between the similarity of the original sources (measured by cosine distance) and success rate in groups containing one hundreds consecutive ShiftNMF simulations each, for sources with the same correlations. Our results demonstrate that for relatively uncorrelated source signals the success rate is above 70%. This rate experiences a sharp drop off to around 30% for signals with greater than 0.6 cosine similarity, and a further decreases to the 10%, when the similarity exceeds 0.8.

3.4 ShiftNMFk reconstructions

By combining ShiftNMF with our NMFk method and with the elimination criteria, we were able successfully to determine the number of sources for a series of synthetic examples discussed below.

3.4.1 Reconstruction of 3 original signals using 18 sensors

We begin with 3 pre-determined signals mixed and delayed randomly to produce a test case with observations recorded by 18 sensors. Figure 2 shows the resulting robustness of the ShiftNMFk solutions. The comparison of the average Silhouette width (in blue) for d from 1 through 6 sources, shows how well our solutions for the original signals are clustered together. The value in the plot is the average of the Silhouettes of each cluster. We observe that there is *no good clustering* for more than 3 sources. The red curve shows how well the average reconstruction error, that is - the difference between the original observations and the reconstruction, is minimized. Again the best result with average reconstruction error closest to zero is achieved for more than 3 sources. We can easily conclude that the system has 3 original sources mixed at 18 sensors. The AIC score, presented in Table 1, confirms this conclusion. The lower plot shows the derived signals and their mixing matrix and respective delays as compared to the true signals, used to generate the example.

3.4.2 Reconstruction of 4 original signals using 24 sensors

Figure 3 shows the results from simulations with 4 sources and 24 sensors. Again, the comparison of the average Silhouette width (in blue) for $d = 1$ through $d = 6$ sources, shows how well our solutions of the original signals are clustered together. The value in the plot is the average of the Silhouettes of each cluster. We see the best clustering occurs for 4 sources. The red curve shows how well the average reconstruction error that is - the difference between the original observations and the reconstruction, is minimized. Clearly, the best result with average reconstruction error closest to zero is achieved for 4 sources. We can easily conclude that the system has 4 original sources mixed at 24 sensors. Again, the AIC score, presented in Table 1, confirms this conclusion. The lower plot shows the derived signals and their mixing matrix and respective delays as compared to the true values used to generate the example.

It is important to note that a combination of the Silhouette width cut-off and the AIC criteria can be applied to select the optimal number of sources. Indeed, based on the results presented in Table 1 and the figures in this section, we observe that applying the Silhouette width cut-off together with the AIC criteria provides reliable estimates of the unknown number of sources. This is important observation considering that both metrics (Silhouette width and AIC) explore very different properties of the estimated solutions. The Silhouette width focuses on the robustness and reproducibility of the solutions, while the AIC score evaluates the solution parsimony.

3.5 Identification of the source locations

To demonstrate the capabilities of our algorithm to identify source location, we generated two additional synthetic examples. We placed (arbitrary) 3 sources: a) inside and b) outside

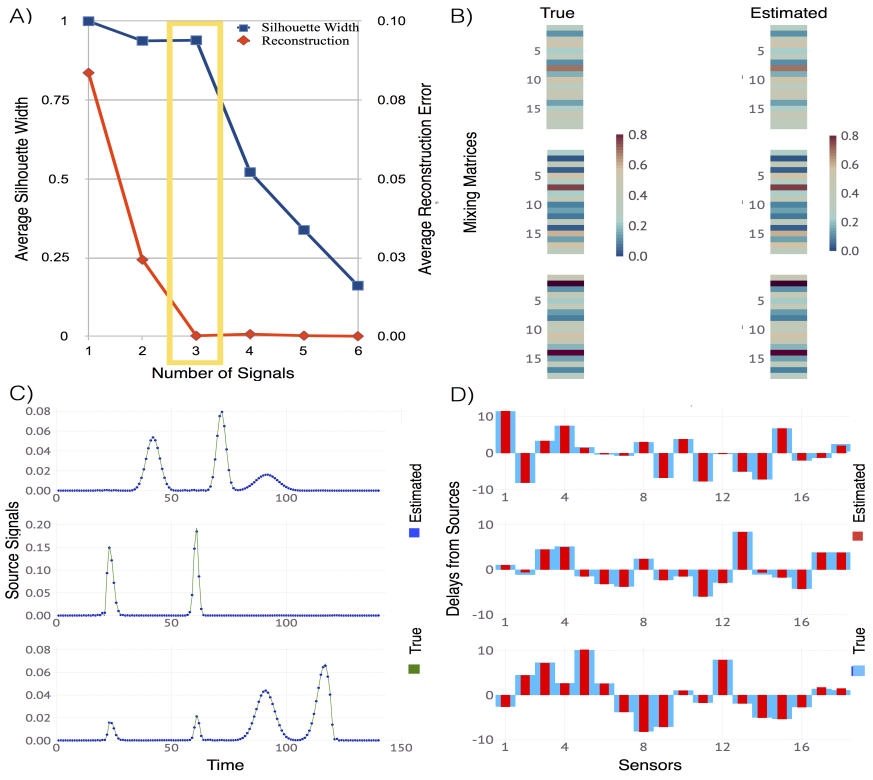

Figure 2: ShiftNMFk algorithm correctly identifying the number of sources, signals, delays and mixing weights. Panel A: Average Silhouette width (blue) and Average Reconstruction Error (red) for each guess of number of sources in the system. Panel B: Color coated visualization of the mixing matrices (true and estimated) where each column represents the weight of a source to each sensor. Panel C: The original (green line) and found (blue dots) signal transients superimposed over each other. Panel D: The delays of the sources to each sensor represented as a bar graph with original(blue bars) and found delays(red bars).

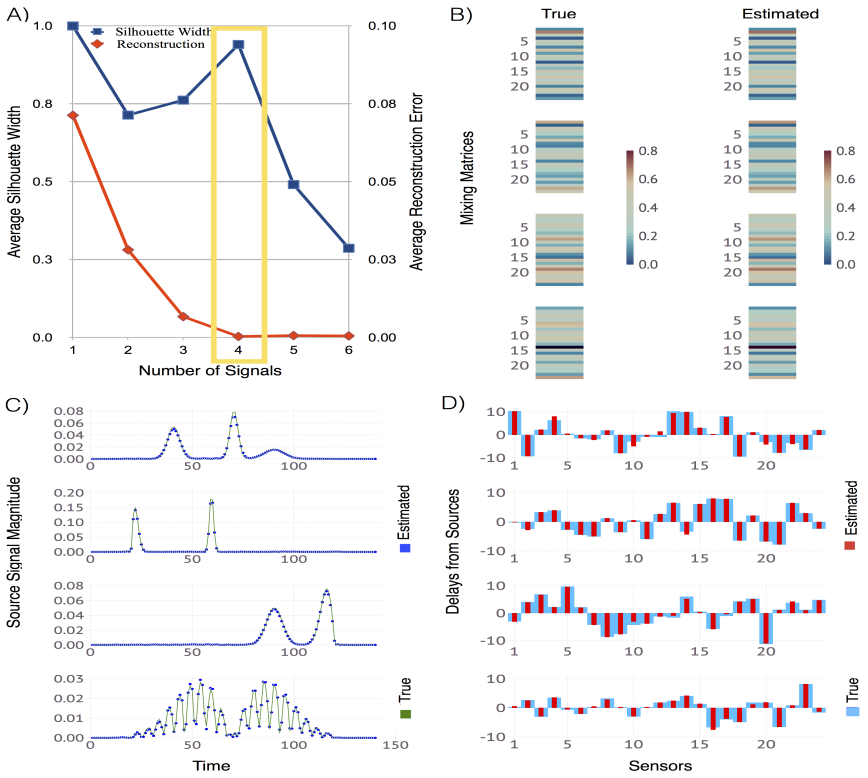

Figure 3: ShiftNMFk results identifying the number of sources, signals, mixing weights, and signal delays. Panel A: Average Silhouette width (blue) and Average Reconstruction Error (red) for different number of sources. Panel B: Color coded visualization of the mixing matrices (true and estimated) where each column represents the weight of a source to each sensor. Panel C: The original (green line) and estimated (blue dots) signal transients superimposed over each other. Panel D: The delays of the sources to each sensor represented as bar graphs with original (blue bars) and found delays (red bars).

Case	$AIC \times 10^{-3}$ d (number of sources)					
	1	2	3	4	5	6
3 sources 18 sensors	-14.043	-15.818	**-22.412**	(-20.789)	(-22.194)	(-23.463)
4 sources 24 sensors	-19.585	-21.367	-24.114	**-29.734**	(-28.597)	(-28.916)
3 sources inside 16 sensors	-13.411	-16.741	**-21.078**	(-18.988)	(-20.260)	-
3 sources outside 16 sensors	-13.383	-16.888	**-21.400**	(-19.232)	(-19.819)	-

Table 1: AIC calculated for each number of original sources d, given its reconstruction error. Its value is minimized by the optimal number of sources. The values of d that are rejected because of bad clustering (Silhouette width < 0.7) are given in parentheses. Note that the Silhouette width cut-off and the AIC criteria have to be combined to select the correct number of sources.

of a grid with 16 evenly spaced sensors. We choose, the waveforms of the original unknown signals to be same as these shown in Figure 2. For these examples, instead of generating random weight, W, and delay, τ, matrices, we calculated them as a function of the arbitrary positions of the sources, that is, of the resulting random distances between each sensor and each source. For simplicity the speed (same for all signals) of the signals in the medium in these examples is set to 1, and we consider the decay of the amplitude of the signals to follow either $1/r$, $1/r^2$ or $1/\sqrt{r}$ law; the specified type is reflected on the values of the mixing matrix elements. These particular types of decay were considered because of their significance in common physical problems. Thus, the $1/r$ and $1/r^2$ power laws reflect the physics of wave propagation in a three-dimensional space – they describe the decay of the amplitude and the intensity (which goes as the square of the amplitude) of a wave with the distance from the source. The last one ($1/\sqrt{r}$) is characteristic of the decay of the amplitude of two-dimensional surface waves in three-dimensional space (e.g., Rayleigh seismic waves). Note that in these examples the signals keep their original waveforms as they propagate in space, i.e. the waves are considered to be dispersionless.

Taking into account one of these types of decays leads also to a specific constraint on the elements of the sought W matrix that can be used to significantly increase the speed and efficiency of the algorithm for determining the source position. For concreteness, we would consider surface waves, but similar reasoning can be applied to other types of waves. If we make the additional assumption that all sensors are identical, we can write for the weight of the j-th source at the i_1-th sensor $w_{j,i_1} = w_0/\sqrt{R_{j,i_1}}$, where w_0 is a constant, which is the same for all sensors, and R_{ji_1} is the source-sensor distance. This general relation is applied to construct upper boundaries for the distances between sources and sensors. We proceed by first forming the ratio of the amplitudes of the signals from the same source at two different sensors, say i_1 and i_2, and using the vector equation $\mathbf{R}_{j,i_2} = \mathbf{R}_{j,i_1} + \mathbf{r}_{1,2}$, where

$\mathbf{r}_{1,2}$ is the vector connecting the two sensors, i_1 and i_2, we get

$$\frac{w_{j,i_2}}{w_{j,i_1}} = \left(\frac{R_{j,i_1}}{R_{j,i_2}}\right)^{1/2} = \frac{\left(R_{j,i_2}^2 - 2\cos\theta_{1,2} R_{j,i_2} r_{1,2} + r_{1,2}^2\right)^{1/4}}{R_{j,i_2}^{1/2}},$$

with $\theta_{1,2}$ being the angle between $\mathbf{R}_{j,i_2}$ and $\mathbf{r}_{1,2}$. Since we want to obtain only an upper bound on the distance between the source and the sensors we can expand the square root, assuming that $r_{1,2} \ll R_{j,i_2}$ (valid if the distance of the source to the array is much larger compared than the size of the array), which gives,

$$\frac{w_{j,i_2}}{w_{j,i_1}} \approx 1 - \frac{r_{1,2}\cos\theta_{1,2}}{2R_{j,i_2}}$$

This finally leads to the result

$$R_{j,i_2} \leqslant \frac{r_{1,2}}{2\left(1 - \frac{w_{j,i_2}}{w_{j,i_1}}\right)}. \tag{21}$$

This formula gives an estimate for the upper boundary of the distance between the source j and the entire sensor array, provided that $w_{j,i_2}/w_{j,i_2}$ is less than 1. Indeed, for a fixed j, w_{j,i_2} and w_{j,i_2} can be chosen to be such as the $w_{j,i_2}/w_{j,i_2}$ to be the smallest possible, e.g., if the selected sensors are the two furthest apart along the direction of the signal propagation. Eq.21 provides just a rough estimate, rather than a strict upper bound. Indeed, to derive it, we used an expansion of $r_{1,2}/R_{j,i_2}$, but if this ratio is not small (i.e., $r_{1,2}/R_{j,i_2} \lesssim 1$) the next order terms can be non-negligible. For example, in the case $\theta_{12} < \arccos(\sqrt{2/3})$, the second-order term is negative, and thus it could push the R_{j,i_2} estimate up. If both of these conditions are satisfied, their combination can lead to the need of a significant upward correction to R_{j,i_2}. As a practical remedy for these potential issues we multiply the upper limit imposed by Eq.21 by a factor of 2, which did not affect the speed and accuracy of our procedure significantly. The results from ShiftNMFk simulations for both setups, with sources inside and sources outside of the sensor array are presented in Figures 4 and 5, respectively. It can be seen that the estimates are very close to the real ones, used to generate these two problems. The estimates for the source coordinates for these two cases (source inside and outside the sensor array) are presented in Figure 6 and their values are showed with the standard deviations obtained by Bayesian analysis (see Appendix A) are given in Table 2.

Table 2 provides summary of ShiftNMFk results for both synthetic examples. The locations estimated by ShiftNMFk are given with their standard deviation obtained by Bayesian Analysis (See Appendix A).

4. Conclusion

We have developed a new methodology, called ShiftNMFk, for blind source separation of signals observed with temporal delay at a series of sensors. The methodology is based on our previous work on NMFk (Alexandrov and Vesselinov, 2014) and on the shiftNMF (Shifted Non-negative Matrix Factorization) algorithm reported by (Mørup et al., 2007).

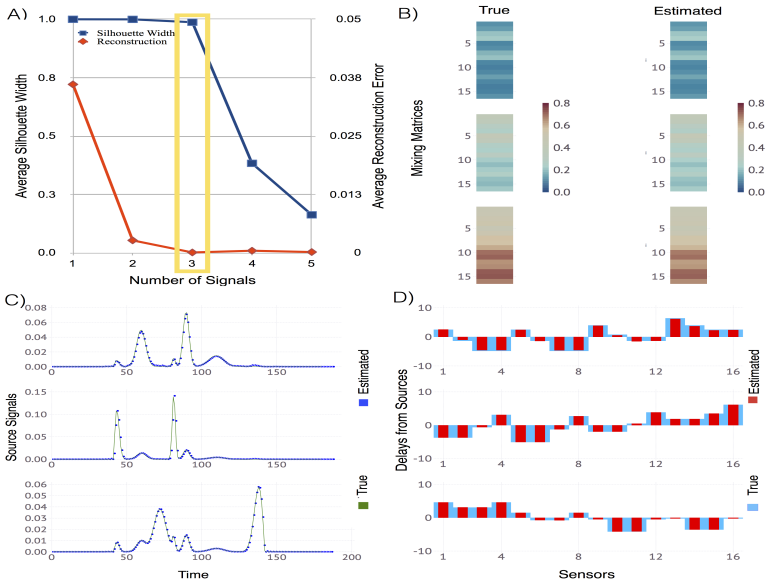

Figure 4: ShiftNMFk results for sources inside the sensor array. Panel A: Average Silhouette width (blue) and Average Reconstruction Error (red) for different number of sources. Panel B: Color coded visualization of the mixing matrices (true and estimated) where each column represents the weight of a source to each sensor. Panel C: The original (green line) and found (blue dots) signal transients superimposed over each other. Panel D: The delays of the sources to each sensor represented as a bar graph with original (blue bars) and found delays (red bars). The correct number of sources is identified as 3 and the transients of the signals as well as their corresponding delays and weights are shown to be very close to the originals.

source position	x_{true}	$x_{NMF} \pm 2\sigma_x$	y_{true}	$y_{NMF} \pm 2\sigma_y$
inside the array	3	3.11 ± 0.07	7	6.95 ± 0.09
	3.5	3.50 ± 2.1e-6	3	2.99 ± 2.9e-6
	6.8	6.80 ± 1.6e-4	5	5.00 ± 5.50e-4
outside the array	-3	-2.98 ± 0.021	6	5.99 ± 2.00e-3
	10	10.01 ± 4.00e-5	3	2.99 ± 6.00e-6
	10.8	10.82 ± 3.0e-4	9.6	9.62 ± 2.35e-4

Table 2: ShiftNMFk estimates for the source coordinates for the two cases depending on the source position (source inside and outside the sensor array) with the standard deviations obtained using Bayesian analysis

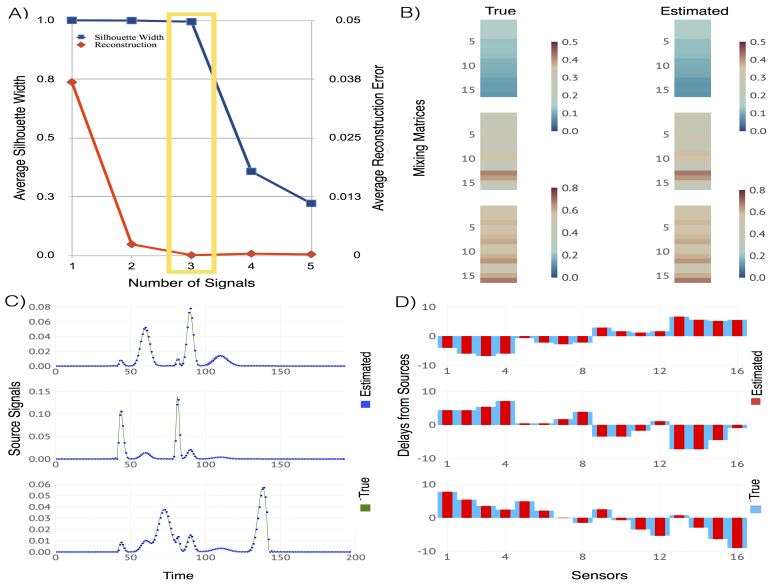

Figure 5: ShiftNMFk results for Sources outside the sensor array. Panel A: Average Silhouette width (blue) and Average Reconstruction Error (red) for different number of sources in the system. Panel B: Color coated visualization of the mixing matrices (true and estimated) where each column represents the weight of a source to each sensor. Panel C: The original (green line) and found (blue dots) signal transients superimposed over each other. Panel D: The delays of the sources to each sensor represented as a bar graph with original (blue bars) and found delays (red bars). The correct number of sources is identified as 3 and the transients of the signals as well as their corresponding delays and weights are shown to be very close to the originals.

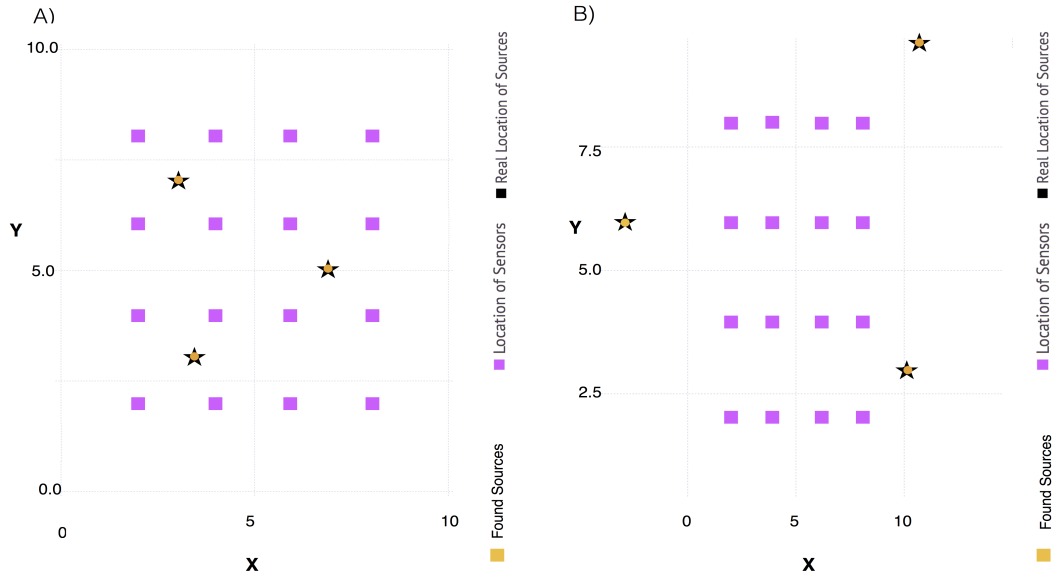

Figure 6: ShiftNMFk results for retrieving the original source location for two cases where the sources are inside (panel A) and outside (panel B) the sensor array.

ShiftNMFk provides an important contribution because the classical NMF does not work when the signals are detected with delays (Mørup et al., 2007), while the original ShiftNMF algorithm requires *a priory* knowledge of the number of the signals. For our method NMFk, we also observed (analyses not shown) that the presence of even very small delays causes substantial problems in the reconstruction of the original signals.

ShiftNMFk analyses presented here demonstrate the applicability of our new methodology for identification of unknown number of delayed sources and their locations and speed, based on the original shiftNMF algorithm combined with (a) our custom semi-supervised clustering based on k-means, (b) a series of selection criteria for rejecting implausible proposal solution, and (c) a custom procedure for identification of source locations. ShiftNMFk also identifies the *velocity of propagation* of each of the signals, based on the signal delays.

For synthetic data sets, the location of the unknown sources are identified from a set of mixed signals recorded by arrays of monitoring sensors, without any additional information about (1) the number of sources, (2) source locations, (3) source propagation velocity, or (4) sources' delays. The solved inverse problem is under-determined (ill-posed). To address this, the ShiftNMFk algorithm explores the plausible inverse solutions and their delays, and seeks to narrow the set of possible solutions by estimating the optimal number of signals needed to robustly, parsimoniously, and accurately characterize the observed data. Based on the estimated signal delays, we can also estimate the source locations and speed. Additional knowledge about the type of the signals and physics of signal propagation (specifically, how the amplitude of the signal changes with the traveled distance) can be leveraged to significantly increase the efficiency and speed-up the convergence of the source-locating procedure.

The original NMFk protocol used only solution robustness based on the Silhouette width to identify the unknown number of sources. Here, we also added a parsimonious criteria based on Akaike Information Criterion (AIC) (Akaike, 2011). The combination of the Silhouette width with AIC was able to select the optimal number of sources. This a very interesting observation considering that the two metrics explore very different properties of the estimated solutions. The Silhouette width focuses on the robustness and reproducibility of the solutions. The AIC scores evaluate the solution parsimony.

An open source Julia language (Bezanson et al., 2012) implementation of ShiftNMFk method, that can be used for identification of a relatively small number of signals, can be found at: https://github.com/rNMF/ShiftNMFk.jl.

The general ShiftNMFk method was submitted as a part of our "Source Identification by Non-Negative Matrix Factorization Combined with Semi-Supervised Clustering," U.S. Provisional Patent App.No.62/381,486, filed by LANL , on 30 August 2016.

Acknowledgments

We would like to express our appreciation to Dr. Ludmil Alexandrov regarding his comments related to the NMFk method. This research was funded by the Environmental Programs Directorate of the Los Alamos National Laboratory. In addition, Velimir Vesselinov was supported by the DiaMonD project (An Integrated Multifaceted Approach to Mathematics at the Interfaces of Data, Models, and Decisions, U.S. Department of Energy Office of Science, Grant #11145687).

Appendix A

Bayesian analysis of posterior uncertainties

Using Bayesian analysis, we can construct regions of likelihood (uncertainties) around the estimated coordinates of the sources. The Bayesian analyses of posterior uncertainties were performed using MADS Vesselinov V. V. et al. (2016). The results are presented in Figures 7 and 8. It can be seen that, the location of a given source and the geometry of the sensors' array could lead to a much narrower uncertainty along some direction, which manifests itself in a non-zero correlation coefficients between x and y coordinates for some sources.

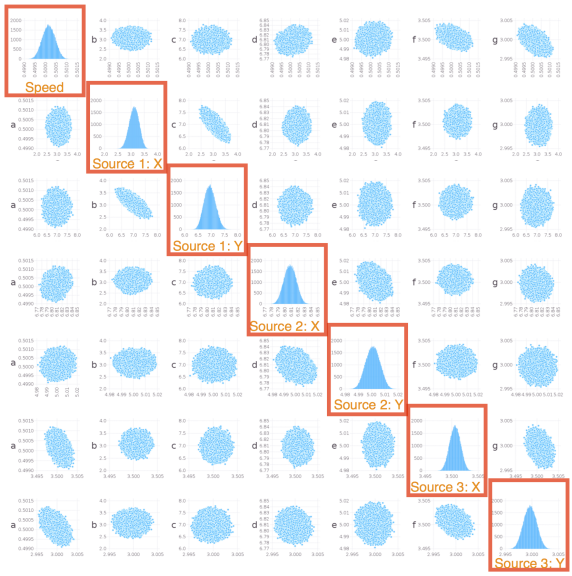

Figure 7: Bayesian analysis of ShiftNMFk results for 3 sources inside of the grid of 16 sensors. The highlighted subplots show the distribution of the variables found by the location minimization procedure; the speed of propagation as well as the x and y locations for each of the 3 sources.

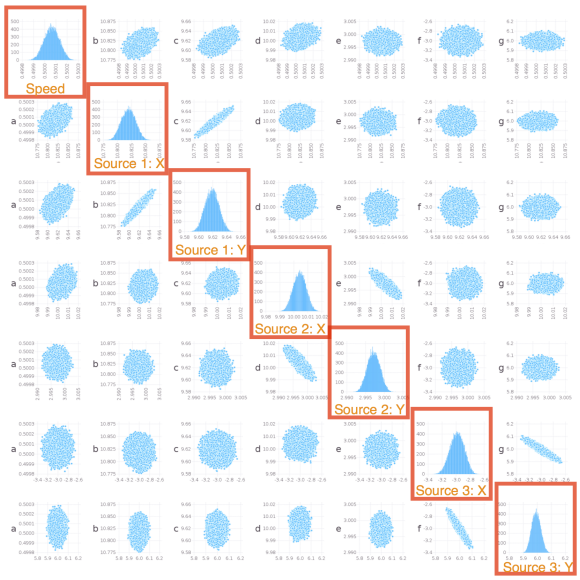

Figure 8: Bayesian analysis of ShiftNMFk results for 3 sources outside a grid of 16 sensors. The highlighted subplots show the distribution of the variables found by the location minimization procedure; the speed of propagation as well as the x and y locations for each of the 3 sources.

References

Hirotugu Akaike. Akaike?s information criterion. In *International Encyclopedia of Statistical Science*, pages 25–25. Springer, 2011.

Boian S Alexandrov and Velimir V Vesselinov. Blind source separation for groundwater pressure analysis based on nonnegative matrix factorization. *Water Resources Research*, 50(9):7332–7347, 2014.

Ludmil B Alexandrov, Serena Nik-Zainal, David C Wedge, Samuel AJR Aparicio, Sam Behjati, Andrew V Biankin, Graham R Bignell, Niccolò Bolli, Ake Borg, Anne-Lise Børresen-Dale, et al. Signatures of mutational processes in human cancer. *Nature*, 500: 415421, 2013a.

Ludmil B Alexandrov, Serena Nik-Zainal, David C Wedge, Peter J Campbell, and Michael R Stratton. Deciphering signatures of mutational processes operative in human cancer. *Cell reports*, 3(1):246–259, 2013b.

Shun-ichi Amari, Andrzej Cichocki, and Howard Hua Yang. A new learning algorithm for blind signal separation. *Advances in neural information processing systems*, 8:757–763, 1996.

Adel Belouchrani, Karim Abed-Meraim, J-F Cardoso, and Eric Moulines. A blind source separation technique using second-order statistics. *IEEE Transactions on signal processing*, 45(2):434–444, 1997.

Jeff Bezanson, Stefan Karpinski, Viral B Shah, and Alan Edelman. Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145*, 2012.

Hsinchun Chen, Roger HL Chiang, and Veda C Storey. Business intelligence and analytics: From big data to big impact. *MIS quarterly*, 36(4):1165–1188, 2012.

Andrzej Cichocki, Rafal Zdunek, Anh Huy Phan, and Shun-ichi Amari. *Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation*. John Wiley & Sons, 2009. ISBN 978-0-470-74666-0.

Martin A Fischler and Robert A Elschlager. The representation and matching of pictorial structures. *Computers, IEEE Transactions on*, 100(1):67–92, 1973.

Jeanny Herault and Christian Jutten. Space or time adaptive signal processing by neural network models. In *Neural networks for computing*, volume 151, pages 206–211. American Institute of Physics Publishing, 1986. doi: 10.1063/1.36258.

Rob Kitchin and Gavin McArdle. What makes big data, big data? exploring the ontological characteristics of 26 datasets. *Big Data & Society*, 3(1):2053951716631130, 2016.

Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

Morten Mørup, Kristoffer H Madsen, and Lars K Hansen. Shifted non-negative matrix factorization. In *Machine Learning for Signal Processing, 2007 IEEE Workshop on*, pages 139–144. IEEE, 2007.

Jeremy E Oakley and Anthony O'Hagan. Probabilistic sensitivity analysis of complex models: a bayesian approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3):751–769, 2004.

Pentti Paatero and Unto Tapper. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2): 111–126, 1994.

Tan Pang-Ning, Michael Steinbach, and Vipin Kumar. *Introduction to data mining*. Addison-Wesley, 2006. ISBN 978-0321321367.

Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

Paris Smaragdis and Judith C Brown. Non-negative matrix factorization for polyphonic music transcription. In *Applications of Signal Processing to Audio and Acoustics, 2003 IEEE Workshop on.*, pages 177–180. IEEE, 2003.

Vesselinov V. V., D O'Malley, Y Lin, Scott Hansen, and Boian Alexandrov. MADS.jl: Model Analyses and Decision Support in Julia, 2016. URL `http://mads.lanl.gov/`.

Matti Vihola. Robust adaptive metropolis algorithm with coerced acceptance rate. *Statistics and Computing*, 22(5):997–1008, 2012.

Robert Weiss. An approach to bayesian sensitivity analysis. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 739–750, 1996.